Privacy in the Autonomous World of Cyber-Physical Systems

Insup Lee and Oleg Sokolsky
PRECISE Center
Department of Computer and Information Science
University of Pennsylvania

1. Motivation and Problem Statement

Privacy concerns in the increasingly interconnected world are receiving much attention from the research community, policymakers, and general public. However, much of the recent and on-going efforts concentrate on privacy in communication and social interactions. The advent of cyber-physical systems, which aim at tight integration between distributed computational intelligence, communication networks, physical world, and human actors, opens new possibilities for developing intelligent systems with new capabilities. Autonomous cars may reduce number of accidents and increase throughputs of transportation networks. Interoperable medical devices can improve patient safety, mitigate caregiver errors, enable personalized treatments, and allow older adults to age in their places. People are increasingly finding themselves in sensor-rich environments that can provide health-related feedback or warn about potential dangers in the environment. At the same time, cyber-physical systems introduce new challenges and concerns about safety, security, and privacy.

Many of the techniques employed by cyber-physical systems require large amounts of data. Advanced capabilities such as autonomy cannot be realized unless sufficient amount of data is available to infer the operating environment of the system and operator's intent. At the same time, this need to know a lot about the human operator and environment increases privacy concerns, reducing public trust in such systems. Both increased capabilities as well as enhanced privacy are beneficial to the society using cyber-physical systems, but it clear that privacy and system capability are not independent of each other and it may not be possible to arbitrarily increase both.

Providing an appropriate level of privacy is complex due to many factors such as personal preferences, cultural differences, economic benefits/incentives, etc. Privacy in CPS requires an integrative approach that involves developing mathematical/logical foundational understanding and frameworks, engineering solutions to support them, and technologies that can be used in applications. It also requires multiple disciplines: economics, theoretical computer science, distributed systems, CPS, robotics, control theory, application domains.

Understanding the resulting trade-offs and their implications for individual users and the society in general raises novel research questions on many levels. Fundamentally, new understanding of privacy in the technological context is needed, along with rigorous mathematical ways of

capturing privacy considerations in their relationship with other design goals for complex cyber-physical systems. We need new mathematical theories that treat privacy in the larger context. Advances in economics of privacy based on multi-variate optimization techniques and game theory are likely to be needed to address some of these fundamental questions. At the technology level, new system development methods and techniques will be needed to make use of the new underlying theory. Furthermore, research on interactions of human actors with cyber-physical systems and, more importantly, interactions between human actors via cyber-physical systems, will be needed to ensure acceptance of the new technologies by the public and policymakers.

2. Scientific and Research Challenges

We believe that the challenges outlined above can be tackled by the layered approach illustrated in Figure 1. The foundation of the project will be formed by fundamental scientific questions that enable but do not depend on specific technological considerations.

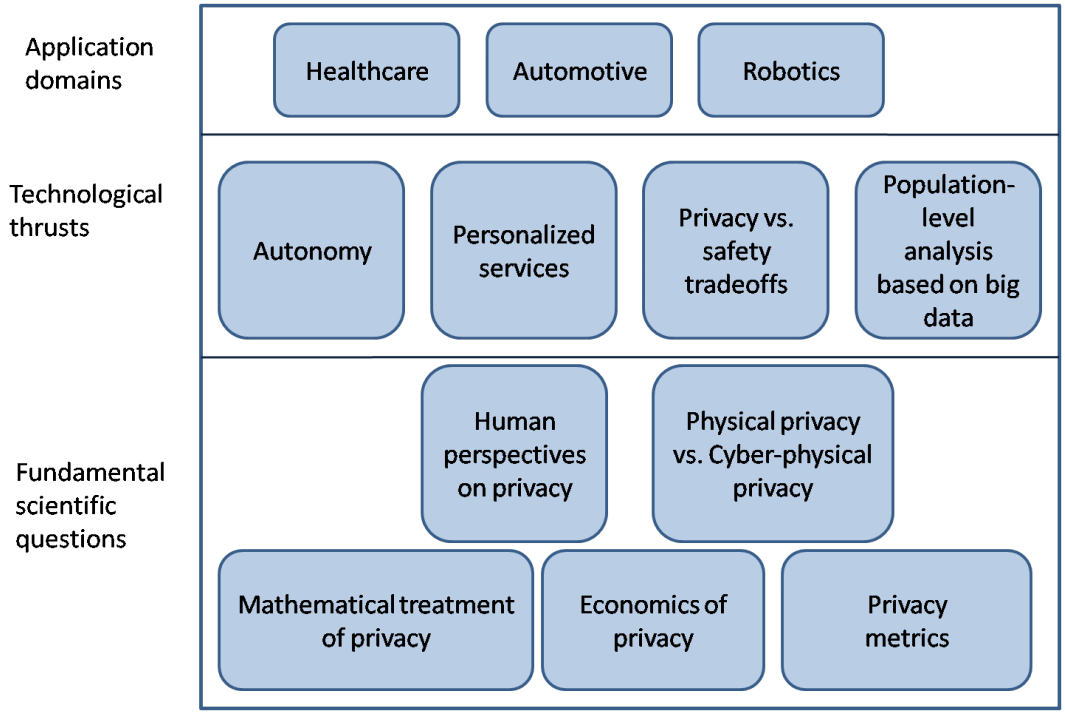

Figure 1. Overview of the proposed approach

2.1. Scientific challenges

At the core, there is an observation that privacy is one of many conflicting considerations in developing modern human-centric systems. Safety, ease of use, and functionality are other

considerations that have to be balanced against each other. The first implication of this observation is that existing mathematical treatment of privacy, such as differential privacy, needs to be modified to allow evaluation of possible tradeoffs.

Furthermore, to enable reasoning about tradeoffs that involve privacy, we need to be able to quantify benefits and risks of exposing data. Doing so requires defining metrics to describe levels of privacy for a given actor or group of actors within a system. Building on these metrics, we can treat the value of privacy in relation to other system aspects that depend on the data provided by actors.

2.2. Technological challenges

Autonomy is one of the central themes in modern cyber-physical systems. Autonomy, that is, ability of the system to achieve a given high-level mission by making localized decision in response to changing circumstances. Key to autonomy is the ability of the system to gain situational awareness based on observing and interpreting changes in its environment. Quality of decisions the system makes depend on diversity and richness of data sources. When some of these sources are restricted due to privacy concerns, the ability of the system to make the right decision can be compromised. We need to develop techniques and tools that would enable system developers to understand how well (if at all) the system would be able to perform a given mission under given privacy constraints.

A related challenge is the interplay between privacy and safety. Many of cyber-physical systems are life critical. In the current system development practice, safety requirements are typically defined independently of privacy considerations. This is an unfortunate situation, because safety considerations determine, how much data the system must have in order to achieve the desired level of safety, and privacy becomes an afterthought. However, if safety and privacy are considered quantitatively together during development, users may settle for less privacy to enhance safety or vice versa, based on personal preferences. It could also allow how to handle various breaking-the-glass scenarios in different emergency response situations..

2.3. Application test beds

To develop privacy-aware technologies and evaluate their effectiveness, there should be test bed for large-scale applications. Medical device domain present numerous opportunities; for example, patient's treatment can be made more efficient and safe when more data about the patient is made available to the system, and older adults can stay in the living environment of their choice longer more safely. Autonomous cars and autonomous robots are also areas of significant expertise that exhibit relevant privacy questions and tradeoffs.